Dr. Faisal Aslam

Tel: +92 42 35723658
Mobile: +92 3324794148
Email: faisal.aslam@gmail.com

Education

Doctorate Of Engineering (Ph.D.) — **October 2006 – March 2011**
University of Freiburg, Germany

M.Sc., Computer Science — **September 2002 – May 2004**
Lahore University of Management Sciences (LUMS)
Pakistan

B.Sc, Computer Science — **December 1997**
Punjab University, Pakistan

Research Interests

- Programming Languages and Software for Embedded Devices
- Protocols and Algorithms for Wireless and Wired Networks
- Traffic Engineering and Network Measurement

Publications

1. **F.Aslam,** L. Fennell, C. Schindelhauer, P. Thiemann, Z. A. Uzmi, "***Rethinking Java Call Stack Design for Tiny Embedded Devices*** ", under review ACM LCTES 2012

2. **F. Aslam**, L. Fennell, C. Schindelhauer, P. Thiemann, Z. A. Uzmi and S. Ruehrup "***Offline GC: Trashing Reachable Objects on Tiny Devices***", 9th ACM Conference on Embedded Networked Sensor Systems (SenSys 2011), Seattle, USA (**19.5% acceptance ratio**)

3. **F. Aslam**, L. Fennell, C. Schindelhauer, P. Thiemann, G. Ernst, E. Haussmann, Z. A. Uzmi and Stefan Ruehrup "***Optimized Java Binary and Virtual Machine for Tiny Motes***", 6th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS 2010), Santa Barbara, USA, 2010.

4. **F. Aslam,** C. Schindelhauer, and A. Vater, "***Improving Geometric Distance Estimation for Sensor Networks and Unit Disk Graphs***". International Workshop on Scalable Ad Hoc and Sensor Networks, (SASN 2009), Saint Petersburg, Russia.

5. **F. Aslam**, C. Schindelhauer, G. Ernst, D. Spyra, J. Meyer and M. Zalloom, "***Introducing TakaTuka - A Java Virtual Machine for Motes***", the 6th ACM Conference on Embedded Networked Sensor Systems, (SenSys 2008, Poster Session), Raleigh, NC, USA.

6. **F. Aslam**, Z. A. Uzmi and A. Farrel, "***Interdomain Path Computation: Challenges and Solutions for Label Switched Networks***", IEEE-Communications Magazine Vol. 45, Issue 10, October 2007 Page(s):94 - 101

7. F. **Aslam**, Z. A. Uzmi, A. Farrel and M. Pióro, "***Inter-domain Path Computation using Improved Crankback Signaling in Label Switched Networks***", IEEE International Conference on Communications 2007 - IEEE-ICC, Glasgow, Scotland, 2007.

8. E. Kubilinskas, **F. Aslam**, M. Dzida and M. Pióro, "***Recovery, Routing and Load Balancing Strategy for an IP/MPLS Network***", 20th International Teletraffic Congress 2007, Ottawa, Canada, 2007.

9. E. Kubilinskas, **F. Aslam**, M. Dzida and M. Pióro, "***Network Protection by Failure-Dependent Backup Paths***", System Science, 2007, Vol. 32, No. 3, pages 5-14

10. **F. Aslam**, S. Raza, and Z. A. Uzmi, "***Bandwidth Sharing with Primary Paths for Protection Routing in an MPLS Network***", Global Internet Workshop of the IEEE INFOCOM, Barcelona, Spain, 2006.

11. E. Kubilinskas, **F. Aslam**, M. Dzida and M. Pióro, "***Network Protection by Single and Failure-Dependent Backup Paths – Modeling and Efficiency Comparison***", System Science, Vol 32, No. 3, Wroclaw University of Technology, Poland, 2006.

12. S. Raza, **F. Aslam**, Shahab Baqai and Z. A. Uzmi, "***Seamless Detection of Link and Node Failures in Label Switched Networks with Local Restoration***", Proc. National Conference on Communications (NCC), New Delhi, India, 2006.

13. S. Raza, **F. Aslam**, and Z. A. Uzmi, "***Online Routing of Bandwidth Guaranteed Paths with Local Restoration using Optimized Aggregate Usage Information***", IEEE International Conference on Communications (IEEE-ICC), Seoul, Korea, 2005.

14. F. **Aslam**, S. Raza, F. R. Dogar, I. U. Ahmad, and Z. A. Uzmi, "***NPP: A Facility Based Computation Framework for Restoration Routing using Aggregate Usage Information***", Third International Workshop on QoS in Multiservice IP Networks (QoS-IP 2005), Catania, Italy, 2005.

Acknowledged In

14. A. Das et al., "***Industrial Strength Ontology Management***", In Proceedings of the First Semantic Semantic Web Working Symposium, SWWS-01, Stanford, USA, August 2001

15. JP. Vasseur, A. Ayyangar, R. Zhang, "*A **Per-domain path computation method for establishing Inter-domain Traffic Engineering (TE) Label Switched Paths (LSPs)***", RFC 5152

Academic Experience

Visiting Assistant Professor
Department of Computer Science
School of Science and Engineering (SSE)
LUMS, Pakistan

1st June 2011 – to Date

Research Associate
Lund University, Sweden.
International Research Support Initiative Program by HEC

December 2005 – May 2006

Lecturer
University of Lahore — **Summer, 2005**
Lahore, Pakistan

Research Assistant
LUMS University Research Program (URP) collaboration — **2003 – 2005**
Cisco Systems, Inc.

Courses Taught

- Topics in Internet Research (grad level) — **Winter 2012, LUMS**
- Computer Networks (under-grad level) — **2005, University of Lahore**

Teaching Assistant (TA)

- Programming Embedded Microsystems, University of Freiburg — **Winter 2009/2010**
- Bachelor Projects in Wireless Sensor Networks, University of Freiburg — **Winter 2008/2009**
- Seminar Ad Hoc Networks, University of Freiburg — **Summer 2008**
- Ad Hoc Networks, University of Freiburg — **Winter 2007/2008**
- Network Protocols and Standards. LUMS, Pakistan — **Fall 2003**

M.Sc. and B.Sc. Thesis Supervised

1. Safdar Iqbal and Abdur Rauf Rajar, "*Extending Offline GC with Path sensitivity and support for multiple threading*", BS thesis, LUMS, 2012
2. S-Y Chao , "*A scalable Wireless Sensor Network Debugger*", M.Sc. Thesis, Uni. Freiburg, October 2010,
3. L. Fennell, "*Garbage Collection for TakaTuka*", B.Sc. thesis, Uni. Freiburg , May 2009
4. M. Zalloom, "*TakaTuka JVM on Sun Spot*", M.Sc. thesis, Uni. Freiburg , October 2008
5. G. M. Ernst, "*Introducing Threads into the TakaTuka Java Virtual Machine*", B.Sc. thesis, Uni. Freiburg, Sept. 2008
6. J. Meyer, "*Garbage Collection for Sensor Motes*", B.Sc. thesis, Uni. Freiburg, July 2008
7. D. Spyra, "*Radio Drivers for Sensor Motes*", B.Sc. thesis, Uni. Freiburg, June 2008
8. A. A. Rahman, "*Implementation of a Massively Parallel Wireless Sensor Network*", M.Sc. Thesis, Uni. Freiburg, 2007

Industry Experience

Verticalnet Inc. Palo Alto, CA, USA — **Feb 2001 – Jan 2002**
Software Engineer

Work on VerticalNet's Ontology Management infrastructure (OMI). For more information: "***Industrial strength ontology management***", In Proceedings of the First Semantic Semantic Web Working Symposium, SWWS-01, Stanford, USA, August 2001

AdventSoft Technologies Inc, Texas, USA — **Jan 2000 – Jan 2001**
Software Engineer and Technical Consultant
www.adventsoft.com

Techlogix Inc, Lahore, Pakistan — **April 1998 – December 1999**
Software Engineer
www.techlogix.com

AdamSoft International and Dedicated IT, Lahore, Pakistan — **March 2002 – June 2002**
Senior Software Engineer
www.adamsoftintl.com

Awards And Honors

- Selected on open merit for *HEC-DAAD* overseas scholarship for Ph.D., 2006.
- Selected for International Research Support Initiative Program by HEC for Lund University, Sweden, 2005.
- Third position in B.Sc. (Computer Science), Punjab University

Technical Skills

Development Languages/Technologies
Java, C, C++, Basic

Database Management Systems and Languages/Tools
Oracle, Access, FoxPro, SQL, PL/SQL, ODBC, DAO, JDBC

Development, Design and Team Management Tools
Net Beans, GCC, IBM Visual Age for Java, Jbuilder 2.0, Visual Café, Visual Source Safe, UML, SVN, CVS

Technical Certifications

IBM Certification in Object Oriented Design with UML — **2001**

Sun Certified Programmer for Java 2 Platform — **2000**

Professional Activities

IEEE-ICC external reviewer — **2007**
Co-Organizer of WEWSN: International Workshop on Energy in Wireless Sensor Networks, co-located with DCOS, Greece — **2008**
External Reviewer for IEEE/ACM Transactions on Networking — **2011**
External Reviewer ACM Transactions on Architecture and Code Optimization — **2011**

3 List Of Publications

List of Publications

1 **F. Aslam**, Ghufran Baig, Mubashir Adnan Qureshi, L. Fennell, C. Schindelhauer, P. Thiemann, Z. A. Uzmi "***Rethinking Java Call Stack Design for Tiny Embedded Devices***", under-review for ACM SIGPLAN LCTES

This paper presents call stack redesign targeted at an efficient use of RAM storage and CPU cycles by a Java program running on a wireless sensor mote. Without affecting the application programs, our call stack redesign saves more than 37% of RAM, on average, evaluated over a large number of benchmarks. On the same set of benchmarks, our design also avoids frequent RAM allocations and deallocations, resulting in about 83% fewer memory operations. These may be critical improvements for tiny embedded devices that are equipped with small amount of RAM and limited battery life. However, our call stack redesign is equally effective for any complex multi-threaded object oriented program developed for desktop computers. We describe the redesign, measure its performance and report the resulting savings in RAM and execution time for a wide variety of programs.

2 **F. Aslam**, L. Fennell, C. Schindelhauer, P. Thiemann, Z. A. Uzmi and S. Ruehrup "***Offline GC: Trashing Reachable Objects on Tiny Devices***" 9th ACM Conference on Embedded Networked Sensor Systems (SenSys 2011), Seattle, USA

The ability of tiny embedded devices to run large, feature-rich Java programs is typically constrained by the amount of memory installed on those devices. Furthermore, the useful operation of such devices in a wireless sensor application is limited by their battery life. We propose Offline Garbage Collection (GC) to alleviate both these limitations. Our approach defies the current practice in which an object may be deallocated only if it is unreachable. The Offline GC allows freeing an object that is still reachable but is guaranteed not to be used again in the program. Furthermore, it may deallocate an object inside a function, a loop or a block where it is last used, even if that object is assigned to a global field. This leads to a larger amount of memory available to a program.
Based on an inter-procedural data flow analysis that is both field-sensitive as well as flow-sensitive, we identify, during program compilation, the point at which an object can safely be deallocated at runtime. We have designed three algorithms for the purpose of making these offline deallocation decisions. To enforce these offline decisions at runtime, the Java bytecode is updated using two customized instructions, also during the compile time.
Our implementation of Offline GC indicates a significant reduction in the amount of RAM and the number of CPU cycles needed to run a variety of benchmark programs. The Offline GC is shown to increase the amount of RAM available to a program by up to 66% compared to a typical online garbage collector. Furthermore, the number of CPU cycles consumed in freeing the memory is reduced by up to 73.13% when the Offline GC is used. From our evaluation of the Offline GC on the benchmarks, we infer that its use is the most effective in saving RAM and CPU cycles for typical sensor network applications that create several dynamic objects.

3 **F. Aslam**, L. Fennell, C. Schindelhauer, P. Thiemann, G. Ernst, E. Haussmann, Z. A. Uzmi and Stefan Ruehrup "***Optimized Java Binary and Virtual Machine for Tiny Motes***", 6th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS 2010), Santa Barbara, USA, 2010

We have developed TakaTuka, a Java Virtual Machine optimized for tiny embedded devices such as wireless sensor motes. TakaTuka requires very little memory and processing power fom the host device. This has been verified by successfully running TakaTuka on four different mote platforms. The focus of this paper is TakaTuka's optimization of program memory usage. In addition, it also gives an overview of TakaTuka's linkage with TinyOS and power management. TakaTuka optimizes storage requirements for the Java classfiles as well as for the JVM interpreter, both of which are expected to be stored on the embedded devices. These optimizations are performed on the desktop computer during the linking phase, before transferring the Java binary and the corresponding JVM interpreter onto a mote and thus without burdening its memory or computation resources.
We have compared TakaTuka with the Sentilla, Darjeeling and Squawk JVMs.

4 **F. Aslam**, C. Schindelhauer, G. Ernst, D. Spyra, J. Meyer and M. Zalloom, "***Introducing TakaTuka - A Java Virtual Machine for Motes***", (Poster Paper) the 6th ACM Conference on Embedded Networked Sensor Systems, (SenSys 2008), Raleigh, NC, USA

We present TakaTuka, a tiny Java Virtual Machine (JVM) for wireless sensor motes. TakaTuka's preliminary version successfully

runs on the mica2 motes using less than 4 KB of RAM. Furthermore, TakaTuka also runs on Windows and Unix.

5 **F. Aslam**, C. Schindelhauer, and A. Vater, "***Improving Geometric Distance Estimation for Sensor Networks and Unit Disk Graphs***". International Workshop on Scalable Ad Hoc and Sensor Networks, (SASN 2009), Saint Petersburg, Russia

Distance measurement between nodes in wireless sensor networks is a prerequisite for a variety of applications and algorithms. However, special hardware allowing such measurements is expensive, especially if dealing with hundreds or thousands of nodes. Fekete et al. presented an approach on distance estimation based on only the neighborhood information available to all nodes in the network. We improve this algorithm, such that it does no longer rely on uniformly distributed nodes. For our approach, it is sufficient that the second derivative of the probability distribution function is a constant.

6 **F. Aslam**, Z. A. Uzmi and A. Farrel, "***Interdomain Path Computation: Challenges and Solutions for Label Switched Networks***", IEEE-Communications Magazine Vol. 45, Issue 10, October 2007 Page(s):94 – 101

For label switched networks, such as MPLS and GMPLS, most existing traffic engineering solutions work in a single routing domain. These solutions do not work when a route from the ingress node to the egress node leaves the routing area or autonomous system of the ingress node. In such cases, the path computation problem becomes complicated because of the unavailability of the complete routing information throughout the network. This is because service providers usually choose not to leak routing information beyond the routing area or AS for scalability constraints and confidentiality concerns. This article serves two purposes. First, it provides a description of the existing and ongoing work in interdomain TE within the IETF. This information is currently found in various Internet drafts and has not yet been collectively presented in a single document. To this end, a summary of both existing path computation architectures - PCE-based and per-domain - is provided. Second, it compares two per-domain path computation schemes in terms of the total number of LSPs successfully placed and average number of domains crossed, without assuming availability of complete topology information. We notice that the two per-domain path computation schemes, proposed in [1, 2], have comparable path computation complexities and setup latencies.

7 **F. Aslam**, Z. A. Uzmi, A. Farrel and M. Pióro, "***Inter-domain Path Computation using Improved Crankback Signaling in Label Switched Networks***", IEEE International Conference on Communications 2007 - ICC-2007, Glasgow, Scotland, 2007

For label switched networks, such as MPLS and GMPLS, most existing traffic engineering (TE) solutions work in a single routing domain. These solutions do not work when a route from the ingress node to the egress node leaves the routing area or the autonomous system (AS) of the ingress node. In such cases, the path computation problem becomes complicated because of the unavailability of the complete routing information throughout the network. We present CWS (computation while switching), a new inter-domain path computation scheme which tries to compute a near-optimal path without assuming the availability of complete topology information. We provide a detailed comparison of the CWS scheme with another per-domain path computation scheme given in J.-P. Vasseur et al. (2006). Unlike the standard per-domain path computation scheme (Vasseur et al., 2006), the CWS scheme continues the quest for a better path instead of terminating the search at the first available path, resulting a significant improvement in terms of path optimality. In particular, CWS guarantees that, for a given network state, a computed inter-domain path will traverse a minimum number of domains. This improvement in path computation directly impacts the amount of traffic that can be allowed on the network. For example, for the COST266 topology with 28 domains and 37 bidirectional inter-domain links, CWS places 960 of the requested 2000 paths as compared to 683 paths placed by existing schemes. Finally, the path setup latency of the CWS scheme remains comparable to that of existing schemes, by allowing the data flow as soon as the first feasible path is found

8 E. Kubilinskas, **F. Aslam**, M. Dzida and M. Pióro, "***Recovery, Routing and Load Balancing Strategy for an IP/MPLS Network***", 20th International Teletraffic Congress 2007, Ottawa, Canada, 2007

The paper considers a problem of routing, protection and load balancing in the IP/MPLS network. A network design problem combining all these aspects is presented. Proportionally fair distribution of residual bandwidths on links is used for load balancing, and protection is achieved with failure–dependent backup paths. The efficiency of the proposed approach is tested by combining optimization and simulation tools. Numerical experiments show that using the proposed load balancing and protection mechanisms decreases the number of disrupted LSPs in case of failures, as compared to other recovery options considered.

9 E. Kubilinskas, **F. Aslam**, M. Dzida and M. Pióro, "***Network Protection by Failure-Dependent Backup Paths***", System Science, 2007, Vol. 32, No. 3, pages 5-14

10 **F. Aslam**, S. Raza, and Z. A. Uzmi, "***Bandwidth Sharing with Primary Paths for Protection Routing in an MPLS Network***", Global Internet Workshop of the IEEE INFOCOM, Barcelona, Spain, 2006

In label-switched networks such as MPLS, protection routing involves computing and setting up the backup paths at the same time when the primary paths are routed. It has previously been shown that two or more backup paths may share bandwidth along common links if such backup paths will never be activated simultaneously. Such sharing between the backup paths leads to reduced bandwidth reservations and, hence improved performance in terms of number of path requests that can be accommodated on the network [1], [2]. We present a novel idea that backup paths may also share bandwidth with certain primary paths, thereby further reducing the overall bandwidth reservations on the network. This results in even more path requests being accommodated on the network. Sharing with primary paths is possible with any protection routing framework. To demonstrate this sharing, we use the NPP protection routing framework as an example [1]. We provide the enhancements to the NPP framework needed to exploit sharing with the primary paths. For the enhanced NPP framework, simulation results on various networks confirm that sharing with primary paths indeed results in better network utilization. This increased performance is achieved with bounded local state information and without requiring any additional routing or signaling overhead.

11 E. Kubilinskas, **F. Aslam**, M. Dzida and M. Pióro, "***Network Protection by Single and Failure-Dependent Backup Paths – Modeling and Efficiency Comparison***", System Science, Vol 32, No. 3, Wroclaw University of Technology, Poland, 2006.

The paper deals with traffic recovery in IP/MPLS network carrying traffic of different priority classes. Two off-line network design problems are presented employing two different recovery mechanisms – single backup path and failure-dependent backup path protection. Also, two routing strategies are studied in this paper. A network design problem with failure-dependent path protection mechanism combined with the routing strategy implying both load balancing and path length minimization is proposed in this paper. The efficiency of the approach proposed is tested by combining optimization and simulation tools. Numerical experiments show that using the proposed protection mechanism allows more traffic to be accommodated in the network and decreases the number of disrupted LSPs in case of failure, as compared to other recovery options considered.

12 S. Raza, **F. Aslam**, Shahab Baqai and Z. A. Uzmi, "***Seamless Detection of Link and Node Failures in Label Switched Networks with Local Restoration***", Proc. National Conference on Communications (NCC), New Delhi, India, 2006.

We consider a label switched network in which backup paths are provided using a local restoration scheme [1], [2]. A node receives a failure notification if either an adjacent link or a neighboring node fails. Such a node, however, cannot distinguish between link and node failures and must do one of two things. 1) Activate backup paths corresponding to both the link and the node suspected to have failed, or 2) Employ a mechanism to establish whether it is a link or a node that has failed, and activate the requisite backup paths once the failure event has been identified. Incase backup paths are activated without disambiguating between link and node failure, bandwidth sharing estimates for a backup path must be revised to take into account concurrent activation of certain additional backup paths. Consequently, a greater amount of bandwidth has to be provisioned when the failure detection is ambiguous. In the case wherein a node waits to identify the exact failure before activating the requisite backup paths, there is increased switchover latency subsequent to the network failure. Evidently, the increased switchover latency translates into a greater traffic disruption following a network failure. We present a simple solution to this problem. Our solution eliminates the need to over-provision backup bandwidth. Moreover, our solution makes possible immediate activation of backup paths without waiting to disambiguate between link and node failure. The key idea is that if an intermediate node along the activated backup paths encounters a resource reservation violation, it can infer the exact type of failure that has transpired. It may then abort the traffic corresponding to the erroneously activated paths while the network traffic that requires restoration can flow through without being disrupted.

13 S. Raza, **F. Aslam**, and Z. A. Uzmi, "***Online Routing of Bandwidth Guaranteed Paths with Local Restoration using Optimized Aggregate Usage Information***", EEE International Conference on Communications (IEEE-ICC), Seoul, Korea, 2005

We investigate the problem of distributed online routing of bandwidth guaranteed paths with local restoration. A unified model is proposed that captures the bandwidth sharing characteristic of backup paths that provision local restoration, corresponding to different fault models. We apply the model to describe bandwidth sharing on backup paths for varying degrees of network state information. The extent of backup bandwidth sharing depends on the amount of network state information made available through routing protocols. A key design criterion for traffic engineering schemes is to maximize the sharing between backup paths, while minimizing this protocol overhead. Kodialam et al. demonstrated in [3] that propagating a constant amount of aggregated information per link leads to cost effective bandwidth sharing. We propose a new optimized aggregate information scenario (oAIS), in which we judiciously select the propagated information, such that the protocol overhead is identical to that in [3].

Simulations show that oAIS outperforms other information scenarios with comparable protocol overhead.

14 F. Aslam, S. Raza, F. R. Dogar, I. U. Ahmad, and Z. A. Uzmi, "***NPP: A Facility Based Computation Framework for Restoration Routing using Aggregate Usage Information***", Third International Workshop on QoS in Multiservice IP Networks (QoS-IP 2005), Catania, Italy, 2005

We present NPP—a new framework for online routing of bandwidth guaranteed paths with local restoration. NPP relies on the propagation of only aggregate link usage information [2,9] through routing protocols. The key advantage of NPP is that it delivers the bandwidth sharing performance achieved by propagating complete per path link usage information [9], while incurring significantly reduced routing protocol overhead. We specify precise implementation models for the restoration routing frameworks presented in [1] and [2] and compare their traffic placement characteristics with those of NPP. Simulation results show that NPP performs significantly better in terms of number of LSPs accepted and total bandwidth placed on the network. For 1000 randomly selected LSP requests on a 20-node homogenous ISP network [8], NPP accepts 775 requests on average compared to 573 requests accepted by the framework of [2] and 693 requests accepted by the framework of [1]. Experiments with different sets of LSP requests and on other networks indicate that NPP results in similar performance gains